BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2014-14

Uncertain Costs and Vertical Differentiation in an Insurance Duopoly

by Radoslav S. Raykov

www.bank-banque-canada.ca

Bank of Canada Working Paper 2014-14

April 2014

Uncertain Costs and Vertical Differentiation in an Insurance Duopoly

by

Radoslav S. Raykov

Financial Stability Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
RRaykov@bankofcanada.ca

Bank of Canada working papers are theoretical or empirical works-in-progress on subjects in economics and finance. The views expressed in this paper are those of the author. No responsibility for them should be attributed to the Bank of Canada.

ISSN 1701-9397

© 2014 Bank of Canada

Acknowledgements

The author would like to thank Uzi Segal, Hanna Halaburda, Hector Perez Saiz, Yaz Terajima, Brian Peterson, and audiences at the Bank of Canada, the IIOC 2013 conference in Boston, and Boston College for beneficial comments and suggestions.

Abstract

Classical oligopoly models predict that firms differentiate vertically as a way of softening price competition, but some metrics suggest very little quality differentiation in the U.S. auto insurance market. I explain this phenomenon using the fact that risk-averse insurance companies with uncertain costs face incentives to converge to a homogeneous quality. Quality changes are capable of boosting as well as reducing profits, since quality differentiation softens price competition, but also undermines the lower-end firm's ability to charge the markup commanded by risk aversion. This can make differentiation suboptimal, leading to a homogeneous quality; the outcome depends on consumers' quality tastes and on how costly quality is. Additional trade-offs between quality costs, profits and profit variances compound this effect, resulting in equilibria at very low quality levels. I argue that this provides one explanation of how insurer competition drove quality down in the nineteenth-century U.S. market for fire insurance.

JEL classification: G22, D43, L22, D81
Bank classification: Market structure and pricing; Economic models

Résumé

Les modèles d'oligopole classiques montrent que les entreprises atténuent la concurrence qu'elles se font sur les prix par la différenciation verticale des produits. Or, certaines mesures laissent croire que le marché de l'assurance automobile aux États-Unis présente un très faible degré de différenciation qualitative. L'auteur explique ce phénomène par le fait que, étant réfractaires au risque et confrontées à des coûts incertains, les compagnies d'assurance sont portées à opter pour des produits de qualité uniforme. L'introduction de différences de qualité entre les produits peut tout aussi bien accroître les profits que les réduire. En effet, si elle rend moins âpre la concurrence sur les prix, la différenciation qualitative limite également la capacité de la firme offrant la variante de moindre qualité d'inclure dans son prix le taux de marge que commande son aversion pour le risque. Selon les goûts des consommateurs pour la qualité et le prix à payer pour l'obtenir, cela peut faire de la différenciation une stratégie sous-optimale qui serait rejetée au profit de l'homogénéisation de la qualité. D'autres arbitrages entre les coûts rattachés à la qualité du produit, les profits et les variances des profits exacerbent cet effet, de sorte que les équilibres se réalisent à de très bas niveaux de qualité. L'auteur soutient que ce serait là une explication au déclin de la qualité observé au XIX^{e} siècle à la suite de l'intensification de la concurrence sur le marché américain de l'assurance incendie.

Classification JEL : G22, D43, L22, D81
Classification de la Banque : Structure de marché et fixation des prix; Modèles économiques

1 Introduction

Classical oligopoly models predict that firms differentiate product quality in order to soften price competition (Shaked and Sutton, 1982; Tirole, 1988). The conventional wisdom is that by making their products distinct, firms better capture consumers with different tastes and extract a larger surplus.

However, in some oligopoly markets, quality differences are muted. One such example is the U.S. market for auto insurance, where auto policies look very similar across sellers in different states, regardless of varying state coverage minima. Insurance markets also differ from conventional markets in two other respects: insurers are risk averse (Mayers and Smith, 1982), and face uncertainty not only about competitor costs, but about their own costs as well. In this paper, I propose a model that accounts for non-differentiated insurance equilibria by showing that they are generated by competition between risk-averse firms with uncertain marginal costs.

The core intuition is compelling, yet has been overlooked by the existing literature. Risk aversion over the cost realization commands firms to charge a positive markup above marginal cost, but the model suggests that the markup is ambiguously affected by quality choice. As a result, quality changes can boost as well as reduce profits: on the one hand, differentiating softens price competition, but on the other, it undermines the low-end firm's ability to charge the markup required by risk aversion. This relationship is key to understanding the main factors behind homogeneous quality.

In a classical differentiated environment, firms normally split the market into price-quality segments to capture clients with different tastes. In a duopoly, therefore, one of the firms pursues a low-price, low-quality strategy, catering to those clients who care less about quality and are willing to pay less. Charging a low price, however, creates a tension with risk aversion. Risk-averse firms require compensation in order to hold risk, and therefore must charge a positive markup above the average marginal cost. If the price required to attract clients with low willingness to pay results in a markup lower than

what is commanded by risk aversion, then differentiating does not make economic sense: the firm incurs a utility loss. Alternatively, if the low-end firm were to charge a higher price for the same quality, it would risk sacrificing more profit from losing demand to its high-end competitor than it could gain from charging the high markup. For this reason, the lower-end firm's profits respond ambiguously to quality changes; this non-linearity can make differentiation unprofitable and enable duopolists to converge in quality.

I find that two major factors influencing quality convergence are the consumers' tastes for quality, which determine their willingness to pay, and the steepness of quality costs. To disentangle the individual influence of each, I consider consumers with heterogeneously distributed quality tastes, and initially assume away the fact that quality is costly. This isolates the influence of tastes on the equilibrium outcome, and shows that quality tastes matter in that they determine the taste point below which less-picky consumers will go to the low-end firm. This taste point determines how low a price the low-end firm can charge in order to secure demand from low-taste consumers while still maintaining a positive markup; the higher the quality taste, the lower the price. But while low pricing secures demand from the low market segment, it also reduces the markup charged over marginal cost, conflicting with risk aversion. If the critical taste point splitting the market turns out to be sufficiently high, then no price charged by the low-end firm will satisfy risk aversion, and the low-end firm will prefer to go high-end as well, since extra quality is free.

Introducing quality costs complicates the picture, but in a way that strengthens the drive to homogeneous quality. I find that the main difference from the free-quality case is that sufficiently steep quality costs can encourage duopolists to converge to the *lowest*-quality outcome, provided consumer tastes for quality are not overly high. The combination of market power and sufficiently steep quality costs, coupled with consumers with low willingness to pay, steers the duopoly to the worst quality point, as this is the most cost-saving scenario given the consumers' ability to pay. This outcome is of special importance if one is trying to understand some episodes from the history of insurance,

especially U.S. insurer competition from the 1850s-70s.

Before the mid-nineteenth century, most states had heavily protectionist policies that favored in-state insurers and discouraged out-of-state entrants (Cummins and Venard, 2007, ch. 2). However, after the industry was deregulated in the mid-1800s to encourage diversification, competition between incumbents and out-of-state entrants depressed premiums in some markets so much that coverage quality soon became inadequate. The market for fire insurance was affected especially badly, since urban fires used to spread quickly and were extremely costly to insurers.[1] Since the correlation of fire damages was driving up insurers' break-even prices, fire coverage was expensive even when sold at a loss, and it was generally difficult to afford. As a result, when competition between incumbents and out-of-state entrants intensified, it soon sent prices below the actuarially fair value. As summarized by Cummins and Venard (2007, Ch. 2):

> As competition increased, premiums dropped and premium adequacy and its corresponding impact on insurer solvency became a notable concern. Efforts by the industry to cooperate by setting prices at adequate levels largely failed. The Great Chicago Fire of October 1871 and Boston fire of November 1872 ($95 million and $50 million in insured losses respectively) resulted in the failure of a large number of insurers and seemed to provide the industry with concrete evidence that premium rates were inadequate.

In the context of standard oligopoly economics, this outcome appears atypical. Oligopolists typically try to exploit market power in order to attain higher-than-market prices and so extract maximum surplus from consumers (for a survey, see Schmalensee and Willig, 1989); therefore, outcomes where an oligopoly market is dominated by low quality at low prices appear unusual. Yet, nineteenth-century fire insurance provides an example where oligopolists deliberately and openly attempted price-setting (even with encouragement from authorities), and nonetheless failed to coordinate. This strongly suggests that high-premium, high-quality insurance was not an equilibrium in this market. One of the contributions of this paper is that it provides an explanation for the

[1] An iconic event in the history of U.S. insurance, the New York fire of 1835 destroyed between 530 and 700 buildings, all insured with New York insurers, some of whom lost their own buildings.

homogeneous low-price – low-quality equilibrium observed historically. This perspective also provides a better understanding of the evolution of U.S. insurance regulation, and suggests why present-day state insurance requirements impose such high minimum standards (discussed in more detail in section 5.2). The idea that high state minimum requirements evolved mainly in order to rule out bad equilibria is entirely supported by the example.

In the rest of the paper, I proceed as follows. First, I review contemporary evidence of the degree of vertical differentiation in the U.S. auto insurance market, and show that anecdotal data from U.S. insurers suggest a very low degree of vertical differentiation and a relatively high quality. On the other hand, the historical example above suggests that non-differentiated equilibria can also occur at low quality levels. To explain *both* phenomena, I build a stylized model of a risk-averse differentiated duopoly with uncertain costs, and argue that cost uncertainty, coupled with risk aversion, is capable of producing homogeneous-quality equilibira at both high and low quality levels.

My strategy is to solve the model analytically. As shown in section 5.2, despite efforts to keep the model simple, tractability turns out to be limited, and closed-form solutions sometimes do not exist. For this reason, instead of aiming to fully categorize all cases where homogeneous quality is possible – something intractable even within this stylized framework – I analyze the main cases where the model has visibility, and argue that the same driving forces apply more generally, such as in an oligopoly market with more than two firms, or to firms with different utility functions or cost distributions. These cases are analyzed in section 6.

2 Cost Uncertainty and Vertical Differentiation in the U.S. Auto Insurance Market

Uncertain costs arise from a variety of sources – uncertain claims, fluctuating input prices, legal risk or exchange rate uncertainty, to name just a few – but are particularly important for insurance companies, where the firm's solvency directly depends on its ability to

accurately predict losses. Insurance companies exist because the pooling of many *i.i.d.* risks reduces aggregate uncertainty, thereby making aggregate losses predictable and loss sharing feasible. In practice, however, predictability is less than perfect even when the number of insured individuals is large, leading to cost uncertainty. Two main reasons contribute to this: small sample size within insurance coverage categories, and small sample size within loss buckets.

An insurance policy contains multiple categories of coverage. For example, a typical U.S. auto policy has seven coverage categories: bodily injury, property damage, bodily injury from uninsured motorist, medical payments, collision damage, comprehensive risk, and personal legal protection. The premium for each category is estimated separately, so that the client can choose between different combinations. To estimate the premium for a category, actuaries use estimates from the subsample of clients who had the particular accident type. Since only a fraction of customers suffer an accident, and some accident types are less frequent than others, loss estimates for some subcategories are based on small samples to which laws of large numbers need not apply.[2]

An estimation problem also arises from small sample sizes within the buckets partitioning the loss distribution. Losses occur more frequently in some buckets than in others, with fewer data points available for extreme loss values. This implies that estimates from the largest-loss buckets also feature the least precision due to large sampling errors, thus introducing uncertainty into the aggregate loss estimate. As pointed out by Wambach (1999), because of this, insurers rarely rely on point estimates of accident probabilities, but rather use distributions instead (OECD, 2005). Since the uncertainty affects the size of claims rather than the insurer's operational cost, insurance models typically model this as a distribution of a flat marginal cost around a known mean (Polborn, 1998; Wambach, 1999; Hardelin and Lemoyne de Forges, 2012). Firms do not

[2]For example, to estimate the premium for bodily injury caused by an uninsured motorist, the insurer can only look at the subsample of policyholders who sustained injuries from an uninsured driver. Likewise, to estimate the premium for legal protection, the insurer can only consider the smaller subsample of clients who had an accident leading to litigation. The uncertain nature of the legal process generates further uncertainty about the length of time a lawyer is needed, and hence about the total legal cost.

know either their own cost realization or the competitor's cost in advance, so they have to choose their quality levels and price their product before actual costs are realized. Information is available only about the cost distribution, which is known in advance.

Another factor specific to the insurance industry is firm risk aversion (Mayers and Smith, 1982), which can likewise arise from a number of sources, such as costly external finance. Costly borrowing is especially relevant for insurers, since whenever losses exceed benchmarks, insurers routinely resort to the capital markets or excess-of-loss contracts with reinsurers. However, this external financing comes at a cost. Typically, verifying the true financial state of the borrower is costly, allowing lenders to charge interest. Reinsurers likewise charge sizable premia, frequently times in excess of the contract's actuarially fair price (Froot, 1999). According to Froot, Scharfstein and Stein (1993), costly external finance is one of the leading factors causing firm risk aversion. Additional factors such as managers' incentives and the convexity of corporate taxes can further contribute to risk-averse firm behavior (Smith and Stulz, 1985). For example, Stulz (1984) argues that firm risk aversion arises from the risk aversion of managers holding their wealth in company stock. Smith and Stulz (1985) argue that corporate taxes, which are convex in earnings, can also lead to risk-averse firm behavior, as a more volatile earnings stream leads to higher average taxes. Bankruptcy costs and capital market imperfections additionally contribute to risk aversion (Smith and Stulz, 1985). As a result, the treatment of insurance companies as risk averters is by now standard in the insurance literature (see Mayers and Smith, 1982; Polborn, 1998; Wambach, 1999; Raviv, 1979; Eliashberg and Winkler, 1981; Blazenko, 1985, among others).

A third factor specific to insurance markets is the surprisingly low degree of vertical differentiation they display by some metrics – especially the U.S. market for auto insurance. If one chooses the extent of coverage as one metric for vertical differentiation, then in a differentiated environment one would expect to see distinct modes in the empirical distribution of demand: for example, a low mode corresponding to policies offering little coverage at low prices, and a high mode corresponding to policies with

extensive coverage at high prices. Nonetheless, anecdotal data compiled by Allstate[3] suggest that the distribution of auto insurance demand has a single mode corresponding to a nearly identical level of coverage across U.S. states. This feature persists regardless of cross-sectional variation in state minimum coverage requirements. For example, the required minimum coverage for bodily injury is $50,000 in Maine, $25,000 in Vermont, and $20,000 in Massachusetts, but the most common coverage sold in these three states is $100,000. The single-mode phenomenon is also highly persistent across the remaining coverage categories. Although not rich enough to permit a firm-level study, Allstate's data are suggestive of a very low degree of vertical differentiation. Moreover, it is worth emphasizing that this homogeneity of the market outcome is not due to a lack of coverage options to choose from; on the contrary, most insurance companies offer customizable policies where the coverage varies with the price. It is therefore this homogeneity of the *equilibrium* outcome that is of interest: consumers and firms appear to be jointly arriving at a non-differentiated equilibrium, while differentiation is possible in principle.

Naturally, vertical differentiation metrics other than coverage also exist. For example, if one traditionally interprets vertical characteristics as "quality," they could also subsume metrics such as the insurer's default probability, service quality and attention to clients, to name a few. I focus on coverage because other quality metrics are either unobservable to consumers (e.g., insurer's default probability), or observable only *ex post* after an accident (e.g., speed of processing claims, quality of contractors doing repairs, general service quality). In addition to being unavailable at the time of purchase, the latter quality metrics also involve a larger degree of subjectivity; in addition, they cannot drive demand *ex ante* unless consumers already have experience with the particular company, or else rely on a signal, such as the insurer's price (Raykov, 2014) or word-of-mouth reputation (a phenomenon studied extensively elsewhere in the *experience goods* literature). For the purposes of this paper, I therefore use the term "quality" generically, with the understanding that many measures of vertical differentiation exist.

[3]Available in tabular form by state at *http://www.allstate.com/auto-insurance/state-coverages.aspx*.

A natural question in the context of differentiated insurers is that of adverse selection. Firms' inability to tell risky from safe clients can produce equilibria with different prices and levels of coverage (separating equilibria), which could appear similar to the differentiated outcome discussed here. The key difference between the adverse selection literature and this model is that adverse selection has implications about product differentiation mostly *within* a firm, not *across* firms. For example, Rothschild and Stiglitz's (1976) classical paper concludes that each firm will offer a high-priced, high-coverage contract and a low-priced, low-coverage contract; a pooling equilibrium does not exist, but, in equilibrium, all firms offer the *same* range of products. Apart from the plausibility of this result, which is contested by the subsequent literature (Hellwig, 1986, 1987; Cho and Kreps, 1987), Rothschild-Stiglitz's environment features no differentiation *across* firms, as it considers a competitive market with identical firms and, therefore, considerations of oligopoly power – which is what enables cross-firm differentiation in the first place – are entirely absent. Thus, adverse selection models do not answer the main question in which I am interested: how insurers that would otherwise prefer to specialize in a specific market segment find it more beneficial not to specialize. For this reason, while acknowledging the importance of adverse selection, I do not make it a main focus of this paper.

Factors other than cost uncertainly can also affect product differentiation. For example, Bester (1998) shows that uncertainty about the competitor's quality can reduce spatial differentiation. The claim that this paper makes, therefore, is *not* that cost uncertainty is the single cause for reduced vertical differentiation in the insurance market; rather, given the specific characteristics of the insurance industry, cost uncertainty stands out as the most plausible explanation.

3 Model

Traditionally, the literature has addressed the question of price-quality choice using a family of duopoly models derived from the well-known papers of Gabszewicz and Thisse

(1979) and Shaked and Sutton (1982). I build a game-theoretic model in the same tradition where a vertical characteristic, for simplicity dubbed "quality," is determined endogenously as the subgame-perfect Nash equilibrium of a two-stage competition game. In the first stage, each firm invests in quality, pays a quality cost, and chooses a quality level. In the second stage, the risk-averse duopolists play a Bertrand competition game with uncertain costs, parameterized by the quality choices made in the first stage. The subgame-perfect Nash equilibrium quality choice is found by backward induction.

I consider two firms, Firm 1 and Firm 2, who can produce a good of quality s_1 and s_2, respectively, charging prices p_1 and p_2. A feature specific to the insurance market is that policies are priced and sold before the insurer knows the true realization of accident costs. Since the uncertainty affects the payment for each individual policy, rather than the insurer's cost of doing business, insurance models typically model this as a distribution of a flat marginal cost around a known mean (Polborn, 1998; Wambach, 1999; Hardelin and Lemoyne de Forges, 2012). Since firms do not know the true cost realization in advance, they have to choose their quality levels and price their product before costs are realized; they can use information only from the cost distribution, which is known in advance. I reflect this in the model by assuming that firms compete in a two-stage game. In the first stage, firms strategically choose quality, and then, given quality levels, enter a second stage where they compete on price. Price competition is a more natural concept for the insurance industry than quantity competition, because insurance policies are written on an as-needed basis as consumers show up, rather than first produced in a targeted quantity and then sold (Rees, Gravelle and Wambach, 1999; Hardelin and Lemoyne de Forges, 2012). I solve the game by backward induction using the standard concept of subgame-perfect Nash equilibrium in pure strategies. Having selected a quality of coverage, and priced and sold their policies, firms await the final realization of costs when the actual profit or loss is realized.[4] The timing of the events

[4]Since the model is short-run, the firm can incur either a profit or a loss without needing to exit the market.

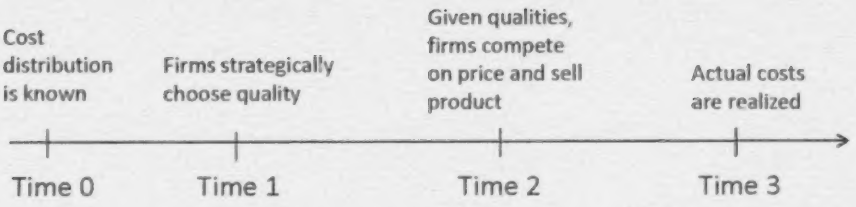

Figure 1: Timing of events in the model

in the model is shown in Figure 1.

Quality can be either costless or costly (see sections 5.1 and 5.2); when costly, the firms pay for it in the first stage, before entering price competition. Quality costs are certain and known in advance; however, the marginal cost is not, because it is driven by accidents expected to occur after the policy has been written. There are no fixed costs.

Quality is a positive real number $s \in [\underline{s}, \overline{s}]$. Since firm indices are arbitrary, we can assume that $s_2 > s_1$ so that the difference in quality is always positive ($\Delta s \equiv s_2 - s_1 > 0$). Due to the mathematics involved, the non-differentiated case ($\Delta s = 0$) is defined as the limiting case of the model when $\Delta s \to 0$. As shown in section 8.2, this limit is well-behaved and is economically meaningful (as differentiation converges to zero, the model converges to classical Bertrand competition with a homogeneous good).

On the demand side, there is a continuum of consumers with heterogeneous tastes and unit aggregate mass, each of whom chooses an insurance provider and can purchase zero or one insurance policy.[5] Consumers care both about the price p and the quality s of the good consumed, summarized by the simple preferences

$$U(s,p) = \begin{cases} \theta s - p & \text{if buys 1 unit} \\ 0 & \text{if no purchase} \end{cases},$$

where θ is a parameter indicating the consumer's quality taste.[6]

Utility is therefore separable in price and quality, and can be interpreted as the surplus

[5] This specification is especially suitable in the context of insurance, where individuals usually purchase at most one policy per property.

[6] A detailed analysis of this preference and its properties appears in Tirole (1988). Mussa and Rosen (1978) consider a variant of the same utility function.

from consuming the good (Tirole, 1988). For a given price, each consumer prefers a higher-quality good to a lower-quality one, but those consumers who have higher values of θ are more willing to pay for high quality (or, equivalently, derive higher surplus); hence θ's interpretation as quality taste. The positive number θ is distributed uniformly over the interval $[\underline{\theta}, \overline{\theta}]$, where $\overline{\theta} = \underline{\theta} + 1$ to ensure unit mass. To economize on notation, I will denote the frequently occurring expression $\overline{\theta}\Delta s$ with $\overline{\Delta}$ and the expression $\underline{\theta}\Delta s$ with $\underline{\Delta}$. To guarantee that the model outcome is non-degenerate, next I impose two assumptions which are standard in the literature (see, e.g., Tirole, 1988):

Assumption 1 (Heterogeneity): $\overline{\theta} > 2\underline{\theta}$.

This assumption makes sure that there is enough consumer heterogeneity to rule out degenerate model behavior (negative equilibrium markups). Notice that, since $\overline{\theta} = \underline{\theta} + 1$, and $\underline{\theta}$ is positive, this assumption restricts the range of $\underline{\theta}$ between $0 < \underline{\theta} < 1$.

Assumption 2 (Market is covered): Minimum quality $\underline{s}$ is large enough that no firm faces zero demand in equilibrium.

Assumption 2 guarantees that the market is a duopoly, as opposed to a monopoly. If minimum quality $\underline{s}$ is so low that at the equilibrium price p_1^* and quality $\underline{s}$ Firm 1's clients prefer to buy nothing (which will happen if $p_1^* > \underline{\theta}\ \underline{s}$), Firm 1 faces the possibility of zero demand, and the model is no longer a duopoly. To prevent this, I require minimum quality $\underline{s}$ to be high enough to guarantee $p_1^* < \underline{\theta}\ \underline{s}$; as will be seen from the expression for the equilibrium price, it is always possible to pick such an $\underline{s}$.

When the market is covered, consumers can choose between a higher- and a lower-quality good sold at two different prices. A consumer of type θ is indifferent between the two goods when

$$\theta s_1 - p_1 = \theta s_2 - p_2.$$

Therefore, given prices and qualities, those consumers with $\theta > \frac{p_2 - p_1}{\Delta s}$ will buy the higher-quality good, and those with $\theta < \frac{p_2 - p_1}{\Delta s}$ will buy the lower-quality good, effectively

splitting the interval $[\underline{\theta}, \overline{\theta}]$ into two parts: those types θ who go to Firm 1, and those who go to Firm 2. Since θ is uniformly distributed over the interval $[\underline{\theta}, \overline{\theta}]$ of length 1, and since quantity demanded is one unit for each person who buys, demand at each firm simply equals the measure (segment length) of consumers positioned on $[\underline{\theta}, \overline{\theta}]$ buying from each respective firm:

$$D_1(p_1, p_2) = \frac{p_2 - p_1}{\Delta s} - \underline{\theta}; \qquad D_2(p_2, p_1) = \overline{\theta} - \frac{p_2 - p_1}{\Delta s}.$$

Each firm's profit equals its quantity demanded times the markup charged above marginal cost c:

$$\Pi_1 = (p_1 - c)\left[\frac{p_2 - p_1}{\Delta s} - \underline{\theta}\right]; \qquad \Pi_2 = (p_2 - c)\left[\overline{\theta} - \frac{p_2 - p_1}{\Delta s}\right].$$

When the marginal cost c is uncertain (reflecting the sampling error in the loss estimate), profits become a random variable with a distribution derived from the distribution of costs. With a large number of insured clients facing independent risks, the distribution of costs is likely to be approximately normal, but for robustness I also consider other distributions (see section 6). Due to the concavity of the firm's utility function, the model's analytical tractability depends on the *joint* choice of utility function and cost distribution. I solve the model for the two cases that are analytically tractable: the CARA[7] utility function with normally distributed costs, and the CRRA[8] utility with uniformly distributed costs; they show that the result is not driven by distributional assumptions, such as the distribution's boundedness or specific shape. This is in line with the core intuition, which does not involve considerations about probability mass. First I present the CARA utility, normal-risk model and use it to illustrate the intuition in detail.

Since, with many *i.i.d.* risks, the marginal cost c will be approximately normally distributed, I initially assume that it follows the normal distribution with mean μ_c and variance σ^2, resulting in normally distributed profits

$$c \sim N(\mu_c, \sigma^2) \implies \Pi \sim N(\mu_\Pi, \sigma_\Pi^2).$$

[7]Constant absolute risk aversion.
[8]Constant relative risk aversion.

The firm's CARA utility function is

$$u(\Pi_k) = -e^{-r\Pi_k} \quad (k = 1, 2),$$

where r is the coefficient of absolute risk aversion. The realization of costs at the two firms is assumed to be independent (firms have no common shocks, since they have no overlapping consumers and accidents are independent).

As shown in Freund (1956), the expected utility of a normally distributed variable $\Pi \sim N(\mu, \sigma^2)$ is

$$\mathbb{E}[-e^{-r\Pi}] = -\exp\left(-r\mu + \frac{1}{2}r^2\sigma^2\right).$$

In the firm's maximization problem, the uncertain profits Π_1 and Π_2 have means and variances as follows:

$$\mu_{\Pi_1} = (p_1 - \mu_c)\frac{p_2 - p_1 - \underline{\Delta}}{\Delta s} \qquad \mu_{\Pi_2} = (p_2 - \mu_c)\frac{\overline{\Delta} - p_2 + p_1}{\Delta s}$$

$$\mathrm{Var}(\Pi_1) = \left[\frac{p_2 - p_1 - \underline{\Delta}}{\Delta s}\right]^2 \sigma^2 \qquad \mathrm{Var}(\Pi_2) = \left[\frac{\overline{\Delta} - p_2 + p_1}{\Delta s}\right]^2 \sigma^2.$$

Substituting mean profits and their variances into the firms' utility functions yields the expected utilities

$$\mathbb{E}u(\Pi_1) = -\exp\left\{-r(p_1 - \mu_c)\frac{p_2 - p_1 - \underline{\Delta}}{\Delta s} + \frac{1}{2}r^2\sigma^2\left[\frac{p_2 - p_1 - \underline{\Delta}}{\Delta s}\right]^2\right\} \tag{1}$$

$$\mathbb{E}u(\Pi_2) = -\exp\left\{-r(p_2 - \mu_c)\frac{\overline{\Delta} - p_2 + p_1}{\Delta s} + \frac{1}{2}r^2\sigma^2\left[\frac{\overline{\Delta} - p_2 + p_1}{\Delta s}\right]^2\right\}, \tag{2}$$

where I will denote the expressions inside the curly braces as ϕ_1 and ϕ_2, respectively.

In the price competition stage, qualities are treated as given and firms compete on price by solving

$$\max_{p_1} \mathbb{E}u(\Pi_1(p, s)) \quad \text{and} \quad \max_{p_2} \mathbb{E}u(\Pi_2(p, s)),$$

where $p = (p_1, p_2)$ is a price vector and $s = (s_1, s_2)$ is a quality vector.

4 Price Competition

In the price competition stage, each firm derives its optimal price reaction to the competitor's price by solving a first-order condition. The fact that the firms' objective functions

(1) and (2) are concave with respect to price is verified in Lemma 1 in the Appendix (section 8.1).

Denoting the expressions in the exponent of equations (1) and (2) as ϕ_1 and ϕ_2, respectively, the first-order condition for Firm 1 can be written as

$$\frac{d\mathbb{E}u(\Pi_1)}{dp_1} = -e^{\phi_1}\frac{d\phi_1}{dp_1} = 0 \iff \frac{d\phi_1}{dp_1} = 0,$$

which, written fully, takes the form

$$-r\left[\frac{p_2 - 2p_1 - \underline{\Delta} + \mu_c}{\Delta s}\right] + r^2\sigma^2\left[\frac{-p_2 + p_1 + \underline{\Delta}}{(\Delta s)^2}\right] = 0.$$

This yields Firm 1's reaction function in terms of p_2:

$$p_1(p_2) = \frac{\mu_c \Delta s + (r\sigma^2 + \Delta s)(p_2 - \underline{\Delta})}{r\sigma^2 + 2\Delta s}.$$

Similarly, Firm 2's first-order condition reduces to

$$-r\left[\frac{\overline{\Delta} - 2p_2 + p_1 + \mu_c}{\Delta s}\right] + r^2\sigma^2\left[\frac{p_2 - p_1 - \overline{\Delta}}{(\Delta s)^2}\right] = 0,$$

resulting in the reaction function

$$p_2(p_1) = \frac{\mu_c \Delta s + (r\sigma^2 + \Delta s)(p_1 + \overline{\Delta})}{r\sigma^2 + 2\Delta s}.$$

Solving the system of reaction functions results in the following unique[9] equilibrium prices:

$$p_1^* = \mu_c + \frac{(r\sigma^2 + \Delta s)[r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s]}{2r\sigma^2 + 3\Delta s}$$
$$p_2^* = \mu_c + \frac{(r\sigma^2 + \Delta s)[r\sigma^2 + (2\overline{\theta} - \underline{\theta})\Delta s]}{2r\sigma^2 + 3\Delta s},$$

where I will often denote the equilibrium markup $(p_k^* - \mu_c)$ as m_k^*.

First observe that, unlike the risk-neutral case, risk-averse firms charge a positive markup above marginal cost, because risk-averse agents require compensation in order

[9]The price competition outcome is unique under this specification, but one can envision specifications where the firm's objective function is not concave with respect to its own price and could in principle result in multiple price equilibria. Unfortunately, those specifications are untractable.

to hold risk (Sandmo, 1971). This softens price competition a bit, but, as will be seen, not enough to reverse the standard results: absent additional incentives, firms will still find it optimal to differentiate maximally even when they are risk averse. As in Tirole (1988), the higher-quality Firm 2 earns higher profits and markup.[10] The Bertrand equilibrium profits and variances are

$$\Pi_1^* = (r\sigma^2 + \Delta s)\left[\frac{r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2, \qquad \Pi_2^* = (r\sigma^2 + \Delta s)\left[\frac{r\sigma^2 + (2\overline{\theta} - \underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2$$

and

$$\mathrm{Var}(\Pi_1^*) = \left[\frac{r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2 \sigma^2, \qquad \mathrm{Var}(\Pi_2^*) = \left[\frac{r\sigma^2 + (2\overline{\theta} - \underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2 \sigma^2.$$

It is intuitive that the higher the firm's risk aversion (or the variance of costs), the higher is the equilibrium markup. Product quality, however, does not always have unambiguous effects on equilibrium profits. The expressions above reveal an interesting trade-off which is essential in the analysis that follows. Whereas Firm 2's profit always increases in own quality, the effect of own quality on Firm 1's profits is generally ambiguous. Nonetheless, above a certain quality level, Firm 1's quality begins to increase the firm's profits, so after some point both firms obtain better profits from higher quality. Higher profits, however, come at the cost of increased variance, as is shown in the next proposition.

Proposition 1. *All else equal,*

(a) Firm 2's profit Π_2^* is increasing in its own quality s_2.

(b) The overall effect of Firm 1's quality s_1 on its profit Π_1^* is ambiguous. However, above some quality s_1^0, Firm 1's quality unambiguously increases its profit $\left(\frac{\partial \Pi_1^*}{\partial s_1} > 0\right)$.

(c) Each firm's equilibrium variance increases in its own quality.

Proof. See the appendix.

[10] This can be shown using the fact that $(2\overline{\theta} - \underline{\theta}) > (\overline{\theta} - 2\underline{\theta})$.

Proposition 1 suggests a trade-off. Above some quality level, higher quality leads to higher expected profits for both firms, but extra profits come at the cost of increased variance. This creates a trade-off between high, variable profits versus lower but more predictable ones. If Firm 1 offers lower quality, it also earns lower profit than Firm 2, but it is very predictable. If Firm 1 offers a higher quality closer to Firm 2, its profits may start growing, but so will their variance. Therefore, forces acting both for and against differentiation exist; the outcome depends on their balance, which is affected by other model parameters (consumers' taste for quality and quality cost).

Substituting the profit and variance equations into the firms' utility functions, the utilities attained in the Bertrand equilibrium are

$$V_1 = \mathbb{E}u(\Pi_1^*) = -\exp\left\{-\left(r\Delta s + \frac{1}{2}r^2\sigma^2\right)\left[\frac{r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2\right\}$$

$$V_2 = \mathbb{E}u(\Pi_2^*) = -\exp\left\{-\left(r\Delta s + \frac{1}{2}r^2\sigma^2\right)\left[\frac{r\sigma^2 + (2\overline{\theta} - \underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2\right\}.$$

At this stage, the firms have computed the prices, profits and utilities that will prevail in the Bertrand equilibrium. Next they find each other's quality reaction functions by maximizing their equilibrium utilities with respect to s.

5 Quality Choice

As I show next, the ambiguous effect of quality on the profits of the low-end firm, combined with the profit-variance trade-off, can sometimes make differentiation suboptimal. The strength of the drive toward non-differentiation will be shown to depend on the consumers' taste for quality. However, other parameters affect the equilibrium as well. For example, quality costs can confound the effect of consumer tastes and obscure their influence on the equilibrium outcome: sufficiently steep quality costs strengthen the drive toward homogeneous quality, although homogeneity is capable of occurring on its own. To isolate these separate driving forces, I first analyze the case where quality is costless, and then consider the added influence of costs. Rather than striving to characterize all conditions resulting in a homogeneous quality (which depend on the model's

specification and tractability), my point is much simpler: to demonstrate by means of a counter-example to the classical model that quality convergence is possible, and that cost uncertainty plays a key role in driving it.

5.1 Costless Quality

In a subgame-perfect Nash equilibrium, each firm knows it has no incentive to deviate from the Bertrand equilibrium prices and quantities, which are parameterized by quality levels. Therefore, when deciding on optimal quality, the firms jointly solve

$$\max_{s_1} \mathbb{E}u(\Pi_1^*(p^*, s), \qquad \max_{s_2} \mathbb{E}u(\Pi_2^*(p^*, s)).$$

Just as in the standard model (Tirole, 1988), when quality is costless, it will turn out that the higher-quality firm (Firm 2) wants to pick the maximum quality regardless of what Firm 1 does; the explanation for this, however, is different. Recall that by Proposition 1, equilibrium variances increase with own quality, so by increasing quality, the firm increases not only its expected profits, but also their variance. The model tells us that in this particular case, the profit effect dominates over that of variance.

Proposition 2. *When quality is costless, Firm 2 will always choose the highest quality,* $\overline{s}$.

Proof. See the appendix.

However, model results diverge from the classical paradigm for the *lower-or-equal-quality* Firm 1. The reason is that the lower-quality duopolist faces a trade-off between maintaining a high-enough markup and keeping its customer base when offering a differentiated product. The next proposition shows that there is a critical quality taste value θ_* above which Firm 1's consumers are so easily attracted by the competitor's higher quality that they require lower prices to remain at Firm 1 than the firm can afford. When this happens, differentiating becomes suboptimal because consumers flow out of the lower-end firm and it loses its market share. (Conversely, when consumer quality

taste is sufficiently low, there is a critical value θ_{**} below which maximum differentiation again becomes optimal.)

Proposition 3. *When quality is costless, and given that the second firm's choice is* $s = \overline{s}$,

(a) When consumer tastes are high enough ($\underline{\theta}$ exceeds some θ_*), Firm 1 will also choose $s = \overline{s}$, *resulting in a* ***non-differentiated*** *equilibrium* $(s_1^*, s_2^*) = (\overline{s}, \overline{s})$ *regardless of the values of* r, σ^2.

(b) Moreover, such a critical value θ_* *always exists.*

(c) When $\underline{\theta}$ *falls below some* θ_{**}, *Firm 1 will choose* $s_1^* = \underline{s}$, *resulting in* ***fully differentiated*** *equilibrium* $(s_1^*, s_2^*) = (\underline{s}, \overline{s})$. *Moreover, such a* $\theta_{**} > 0$ *always exists as long as* $r\sigma^2 > 3$.

Proof. See the appendix.

To understand the mechanics behind the result, first note that Firm 1's equilibrium demand D_1^* and equilibrium markup m_1^* both decrease with consumer quality taste $\underline{\theta}$:

$$\frac{\partial D_1^*}{\partial \underline{\theta}} = -\frac{\Delta s}{2r\sigma^2 + 3\Delta s} < 0; \quad \frac{\partial m_1^*}{\partial \underline{\theta}} = -\frac{(r\sigma^2 + \Delta s)[2r\sigma^2 + 3\Delta s]\Delta s}{(2r\sigma^2 + 3\Delta s)^2} < 0.$$

Therefore, for a fixed quality, a higher-quality taste unambiguously lowers equilibrium profits ($\partial \Pi_1^* / \partial \underline{\theta} < 0$). Therefore, when faced with selective consumers, Firm 1 may try to improve profits by instead offering higher quality in order to capture more consumers to boost demand, because

$$\frac{\partial D_1^*}{\partial s_1} = -\frac{[2(1 - \underline{\theta}) - 3]r\sigma^2}{(2r\sigma^2 + 3\Delta s)^2} > 0.$$

Higher demand alone, however, may not be enough to increase profits (relative to the fully-differentiated case), because profit also depends on the equilibrium markup.

The behavior of the equilibrium markup with respect to quality taste is the decisive factor determining whether the traditional result will reverse, because it is the high

markup that drives Firm 1's price up and thereby stimulates clients to switch. With less-selective consumers, equilibrium markup falls in own quality, stimulating the firm to stay at the lowest-quality point $\underline{s}$; this is how we get the classical, maximum-differentiation equilibrium in Proposition 3 (c). At high-enough levels of $\underline{\theta}$, however, equilibrium markup begins to increase in quality, so now both demand and the markup pull profits in the upper direction. This can be seen from the behavior of the derivative

$$\frac{\partial m_1^*}{\partial s_1} = -\frac{(1-2\underline{\theta})r^2\sigma^4 + (1-\underline{\theta})[3(\Delta s)^2 + 4r\sigma^2\Delta s)}{(2r\sigma^2 + 3\Delta s)^2}.$$

When $\underline{\theta} < 1/2$, higher quality lowers the markup $\left(\frac{\partial m_1^*}{\partial s_1} < 0\right)$, although it increases demand. If we keep lowering the quality taste $\underline{\theta}$, eventually there comes a point θ_{**} where the markup effect prevails over that of demand: at low levels of $\underline{\theta}$ profits are falling in own quality, which stimulates the firm to stay at the left corner $\underline{s}$. The fact that the variance of equilibrium profits increases in own quality only reinforces this result, because Firm 1's profit variance is lowest at $\underline{s}$; this is the intuition behind the traditional result in Proposition 3 (c).

However, when $\underline{\theta} > 1/2$, the numerator is positive and markup begins to *increase* in own quality: $\frac{\partial m_1^*}{\partial s_1} > 0$. Therefore, when quality taste is high, quality increases both the demand and the equilibrium markup, pushing up profits as well. This is the main force behind the drive to non-differentiation in part (a). However, profits alone cannot tell whether high quality is optimal, because quality-driven profits come at the cost of increased variance (see Proposition 1(c)). Based on intuition alone, one cannot predict which profit-variance combination is better. The model resolves this trade-off by showing that when $\underline{\theta} > \theta_*$, Firm 1 will unambiguously shoot for the high-quality, non-differentiated option. This shows that the utility from extra profits dominates the disutility from increased variance.

To summarize the mechanics, the effect of quality on equilibrium demand, markup, profits and profit variances is shown in Table 1 for different values of $\underline{\theta}$. The intuitive explanation behind non-differentiation is Firm 1's inability to charge low-enough markup

Variable	Sign of $\partial(\cdot)/\partial s_1$	
	$\underline{\theta} < \theta_{**}$	$\underline{\theta} > \theta_{*}$
Equilibrium demand	+	+
Equilibrium markup	−	+
Equilibrium profits	−	+
Equilibrium variance	+	+

Table 1: Effect of quality on demand, markup, profits and profit variances in equilibrium

to keep its consumers from switching when they are too sensitive to quality. In addition to this core intuition, however, the table suggests that the trade-off between profits and profit variances also plays an important role and can confound the drive to non-differentiation.

This discussion makes it easier to understand why risk aversion is essential to my results. Without risk aversion, the equilibrium markup in this model is zero, since the second-stage game reverts to the textbook Bertrand model where prices are driven down to mean marginal cost. Risk-averse firms require a strictly positive markup as compensation for holding risk, which enables product quality to affect the equilibrium markup and profits, sometimes making differentiation suboptimal. Additionally, without risk aversion, the trade-off between quality-driven profits and profit variances shown in Proposition 1 would likewise be moot. Therefore, cost uncertainty in a risk-neutral environment is not enough to reverse differentiation: it is the combination of uncertain costs and risk-averse firms that creates a drive toward homogeneous quality.

5.2 Costly Quality

In realistic settings, attaining high quality requires costly investments. For example, U.S. firehouses were initially built and maintained by local insurers, requiring large up-front investments in equipment and infrastructure. In this section I discuss the effect of costly quality on the differentiation outcome by introducing quality costs. Since very few cost specifications result in tractable solutions, I focus on the particular case of convex (quadratic) quality costs, which is sufficient to produce a counter-example to the classic

results of Shaked and Sutton (1982), and establish my main point that cost uncertainty is essential.

Assume that both firms can invest in product quality s at cost $C(s) = bs^2$, where $b > 0$. Since quality costs are paid in the first stage of the game, they reduce the profits obtained in the second-stage equilibrium one-for-one without changing their variance. The equilibrium utilities attained in the Bertrand equilibrium therefore are

$$V_1 = -\exp\left\{-\left(r\Delta s + \frac{1}{2}r^2\sigma^2\right)\left[\frac{r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2 + rbs_1^2\right\}$$

$$V_2 = -\exp\left\{-\left(r\Delta s + \frac{1}{2}r^2\sigma^2\right)\left[\frac{r\sigma^2 + (2\overline{\theta} - \underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2 + rbs_2^2\right\}.$$

Each firm's solution to the problem $\max_{s_k} V_k(s_k|s_{-k})$ $(k = 1, 2)$ provides its reaction function with respect to the competitor's quality. Even with quadratic quality costs, however, the reaction functions have no closed-form expression, so the Nash equilibrium solutions are analytically untractable. Nonetheless, it is still possible to characterize non-differentiated equilibria at the lowest- and highest-quality point.

Proposition 4. *When quality costs for both firms are given by $C(s) = bs^2$,*

*(a) If $\underline{\theta} < 1/2$ and $b > \frac{5}{16s}$, both firms arrive at a non-differentiated equilibrium with **minimum** quality.*

*(b) If $\underline{\theta} > \frac{1}{2} + 8b\overline{s}$ and $b < \frac{1}{16s}$, both firms arrive at a non-differentiated equilibrium with **maximum** quality.*

Proof. See the appendix.

As this example illustrates, sufficiently steep quality costs can amplify the drive toward a homogeneous quality. If consumers have sufficiently low ability to pay and quality costs are steep, then the duopolists can economize on costs by simultaneously *reducing* quality. This is how the market ends up in the minimum-quality equilibrium of Proposition 4(a). In this case, the availability of market power, combined with high quality costs and consumers with low ability to pay, steers the duopoly into the worst-quality scenario.

As discussed in the Introduction, these two characteristics – an oligopoly market with high quality costs and consumers with relatively low paying ability – characterize well the nineteenth-century U.S. fire insurance market. In the early days of insurance, firehouses were built and maintained by insurers themselves, so providing quality fire insurance required costly upfront investments. The higher correlation of fire damage compared to regular property and casualty losses, on the other hand, made coverage expensive for customers even when sold at break-even prices. Thus the fire insurance oligopoly featured both steep quality costs and consumers with limited ability to pay, the two conditions likely to result in a homogeneous low-quality equilibrium.

With limited competition before the deregulation of the 1850s, the market was stable, but once firm entry intensified price competition, prices dropped below the actuarially fair value and coverage quality tumbled down to a point that was clearly unsustainable in the long run. The interesting aspect about this low-price, low-quality point is that it nonetheless appears to have been a *short-run* equilibrium. As argued by Cummins and Venard (2007), "efforts by the industry to cooperate by setting prices at adequate levels largely failed," implying that deviations from the status quo were unprofitable. In the long term, the market stabilized only after a few trigger events (the Great Chicago Fire of October 1871 and the Boston fire of November 1872) inflicted losses so large as to lead to the failure of a large number of insurers.

It appears reasonable to conjecture that the subsequent evolution of minimum regulatory standards for insurance may have served as a tool to rule out precisely such bad equilibria.[11] This perspective is reinforced by the industry's long-held views that excessive competition among oligopolistic insurers is harmful to market stability (Harrington, 2000). Unlike other industries, where minimum quality standards are low or even non-existent, insurance regulation has evolved over time to make minimum stan-

[11]Studies on minimum quality standards (Garella and Petrakis, 2008; Crampes and Hollander, 1995; Ecchia and Lambertini, 1997; Maxwell, 1998; Valletti, 2000; Toshimitsu and Jinji, 2004) explain how such standards affect product differentiation, collusion, and competition once introduced, while this paper provides insight into why they needed to evolve for insurance in the first place.

Quality tastes	Quality marginal cost b	
	Low or 0	High
Low	Differentiation	Homogeneous low quality
High	Homogeneous high quality	Differentiation

Table 2: Influence of quality tastes and quality costs on the equilibrium outcome

dards quite high. For example, the median bodily injury requirement for auto insurance across U.S. states is $25,000, with some states (Alaska, Maine and Wisconsin) going as high as $50,000. It is a reasonable hypothesis that there is something idiosyncratic about the insurance industry to justify such high minimum standards; I argue that cost uncertainty and risk aversion are the two fundamental factors that set insurance apart. As shown here, these two characteristics make both high-quality and low-quality equilibria possible at homogeneous quality levels, which is consistent with both present-day data on auto insurance and with historical data on fire insurance.

Comparing the results for quadratic quality with the results for costless quality from section 5.1, one can systematize the influence of costs and quality tastes on the equilibrium outcome. When the quality cost is sufficiently steep and consumers have less-demanding tastes, the duopolists converge to a low-quality outcome. Duopolists converge to a high-quality outcome when consumers are selective about quality and the quality marginal cost is low or zero; in the remaining cases, they differentiate. This information is summarized in Table 2.

6 Robustness

6.1 Robustness to the Choice of Utility Function and Cost Distribution

As explained in section 3, the model's analytical tractability strongly depends on the joint choice of firm utility functions and the assumed cost distribution. Few choices of a utility function and cost distribution result in a tractable combination. However, the model provides enough visibility for the two most important utility function classes – with constant absolute and relative risk aversion, and the two most relevant risk distributions

– the normal and the uniform. The analysis so far has focused on the combination of CARA utility and a normal distribution of costs. The proposition below extends the analysis to the CRRA utility with a risk-aversion coefficient of $R = 2$ and uniformly distributed costs.[12] It confirms that my main results are not driven by the choice of utility or by specific properties of the cost distribution, such as boundedness or normality. This conforms with the underlying economic intuition: any non-degenerate risk distribution should result in a positive markup, whose ambiguous response to quality is what drives non-differentiation.

Proposition 5. *Let the firms' utility function be given by* $u(x) = \frac{x^{1-R}}{1-R}$ *with* $R = 2$*, and costs be uniformly distributed on an interval* $[\underline{c}, \overline{c}]$ *with mean* μ *and variance* σ^2*, with remaining assumptions unchanged. Then, if quality costs for both firms are given by the strictly increasing function* $C(s)$*, firms arrive at a* ***non-differentiated*** *equilibrium at the* ***minimum*** *quality.*

Proof. See the appendix.

6.2 Oligopoly with More than Two Firms

The model's tractability is likewise limited with more than two firms, but the economic intuition is straightforward enough to assess the multi-firm outcome qualitatively. Multiple firms will want to segment the market into multiple price-quality categories, so that each firm can capture consumers with a particular taste range and so attain a "local monopoly power" over them to extract more surplus. In this environment, low pricing is likely to conflict with the markup required by risk aversion only for firms at the lower end of the price-quality spectrum. Thus, selective quality tastes may make differentiation suboptimal in some of the lower-end, but not all market segments. The result will then likely be a compression of the quality distribution, rather than a drive to a fully homogeneous quality. Likewise, quality costs may be unprofitably steep for the lowest-end but not medium or high-end firms, who depend on quality to maintain their customer base.

[12] Other levels of risk aversion are likewise untractable in the CRRA case.

Again, a compression of the quality distribution is most likely. The remaining degree of differentiation seems consistent with the small but existent differences between auto insurers on dimensions other than coverage – such as service quality, speed of processing claims and overall customer satisfaction, which, albeit subjective, are still valid reasons to choose one insurer over another.

7 Conclusion

Cornerstone oligopoly models predict that firms differentiate quality as a way of softening price competition. Muted quality differences in the market for auto insurance, however, suggest that vertical differentiation is not always optimal. I explain the low degree of quality differentiation in the auto insurance market using the fact that risk-averse insurance companies with uncertain costs face incentives to converge to a homogeneous quality. Risk aversion commands a markup above the average marginal cost, which can undermine the lower-quality firm's ability to maintain low-enough prices to retain its customers, making the low-price, low-quality strategy unprofitable and resulting in non-differentiation. However, an additional trade-off between profits and profit variances compounds this effect and, taken jointly with the effect of quality cost, can produce non-differentiated outcomes at very low quality levels; the outcome depends on the balance between consumer quality tastes and quality cost. I argue that the existence of a homogeneous low-quality equilibrium explains why insurance markets have historically featured high minimum standards compared to other industries, and that this provides a better understanding of the historical evolution of insurance regulation. Historical examples from the nineteenth-century U.S. market for fire insurance confirm the existence of low-price, low-quality equilibria before insurance regulation evolved to include minimum quality standards. While many factors can contribute to a non-differentiated insurance market, uncertain costs coupled with firm risk aversion stand out as a plausible explanation, since it is precisely the combination of risk aversion and cost uncertainty that makes the insurance sector unique.

8 Appendix

8.1 Proofs

Lemma 1. *Both firms' utility funcions are concave with respect to the own price. In particular,*

(a) $\mathbb{E}u(\Pi_1)$ is strictly concave in p_1.

(b) $\mathbb{E}u(\Pi_2)$ is strictly concave in p_2.

Proof. Recall that

$$\mathbb{E}u(\Pi_1) = -\exp\left\{-r(p_1-\mu_c)\frac{p_2-p_1-\underline{\Delta}}{\Delta s}+\frac{1}{2}r^2\sigma^2\left[\frac{p_2-p_1-\underline{\Delta}}{\Delta s}\right]^2\right\} = -e^{\phi_1},$$

$$\mathbb{E}u(\Pi_2) = -\exp\left\{-r(p_2-\mu_c)\frac{\overline{\Delta}-p_2+p_1}{\Delta s}+\frac{1}{2}r^2\sigma^2\left[\frac{\overline{\Delta}-p_2+p_1}{\Delta s}\right]^2\right\} = -e^{\phi_2},$$

labeling the expressions inside the exponents as ϕ_1 and ϕ_2, respectively. I will first show that ϕ_1 is convex in p_1 and therefore $d^2\mathbb{E}u(\Pi_1)/dp_1^2 < 0$.

An inspection of the derivatives of ϕ_1 shows that

$$\frac{d\phi_1}{dp_1} = -r\left[\frac{p_2-2p_1-\underline{\Delta}+\mu_c}{\Delta s}\right]+\sigma^2r^2\frac{p_1-p_2+\underline{\Delta}}{\Delta s}.$$

$$\frac{d^2\phi_1}{dp_1^2} = \frac{2r}{\Delta s}+\sigma^2\left(\frac{r}{\Delta s}\right)^2 > 0.$$

Therefore ϕ_1 is strictly convex in p_1. Taking this into account, we can now look at the sign of

$$\frac{d^2}{dp_1^2}[-e^{\phi_1}] = -e^{\phi_1}\left[\left(\frac{d\phi_1}{dp_1}\right)^2+\frac{d^2\phi_1}{dp_1^2}\right] < 0,$$

which is negative because $-e^{\phi_1} < 0$, while the convexity of ϕ_1 ensures that $\frac{d^2\phi_1}{dp_1^2} > 0$.

The same procedure applied to Firm 2 reveals that ϕ_2 is likewise convex in p_2 and that therefore

$$\frac{d^2}{dp_2^2}[-e^{\phi_2}] = -e^{\phi_2}\left[\left(\frac{d\phi_2}{dp_2}\right)^2+\frac{d^2\phi_2}{dp_2^2}\right] < 0. \quad \blacksquare$$

Proposition 1. *All else equal,*

(a) Firm 2's profit Π_2^* is increasing in its own quality s_2.

(b) The overall effect of Firm 1's quality s_1 on its profit Π_1^* is ambiguous. However, above some quality s_1^0, Firm 1's quality unambiguously increases its profit $\left(\frac{\partial \Pi_1^*}{\partial s_1} > 0\right)$.

(c) Each firm's equilibrium variance increases in its own quality.

Proof. (a) The partial derivative

$$\frac{\partial \Pi_2^*}{\partial s_2} =$$
$$= \frac{4(1+\underline{\theta})r^3\sigma^6 + 3(2+\underline{\theta})^2(\Delta s)^3 + 6r\sigma^2(2+\underline{\theta})^2(\Delta s)^2 + r^2\sigma^2(4\underline{\theta}^2 + 18\underline{\theta} + 17)\Delta s}{(3\Delta s + 2r\sigma^2)^3} > 0$$

is strictly positive because all terms are ≥ 0, while $4(1+\underline{\theta})r^3\sigma^6 > 0$. $\square$

(b) The partial of Π_1^* with respect to s_1 is

$$\frac{\partial \Pi_1^*}{\partial s_1} = -\frac{3[1+2r\sigma^2](1-\underline{\theta})^2(\Delta s)^2 + r^2\sigma^2(4\underline{\theta}^2 - 10\underline{\theta} + 3)\Delta s - 4\underline{\theta}(r\sigma^2)^3}{(3\Delta s + 2r\sigma^2)^3}.$$

In the general case, the sign of $\frac{\partial \Pi_1^*}{\partial s_1}$ is ambiguous because it depends on the values of Δs, $\underline{\theta}$, r and σ^2. However, observe that as $\Delta s \to 0$,

$$\frac{\partial \Pi_1^*}{\partial s_1} \to \frac{4\underline{\theta}(r\sigma^2)^3}{3\Delta s + 2r\sigma^2} > 0.$$

Recalling that $\Delta s \equiv s_2 - s_1$, next observe that $\partial \Pi_1^*/\partial s_1$ is continuous in s_1, which implies that for any given level of s_2, there exists a neighborhood where the function's sign is preserved. Therefore, there is some $s_1^0 < s_2$ above which the sign of $\partial \Pi_1^*/\partial s_1$ remains positive:

$$\text{Given } s_2, \quad \exists s_1^0 : \quad \partial \Pi_1^*/\partial s_1 > 0 \quad \text{for} \quad s_1 > s_1^0. \quad \square$$

(c) If we put

$$\left[\frac{r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2 \equiv g_1(s_1) \qquad \left[\frac{r\sigma^2 + (2\overline{\theta} - \underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2 \equiv g_2(s_2),$$

then we can write $\text{Var}(\Pi_1^*) = g_1(s_1)\sigma^2$ and $\text{Var}(\Pi_2^*) = g_2(s_2)\sigma^2$. To find the effect of qualities on the variances, we need the derivatives g_1' and g_2':

$$g_1'(s_1) = \frac{2r\sigma^2[r\sigma^2 + (1-\underline{\theta})\Delta s][1+2\underline{\theta}]}{[2r\sigma^2 + 3\Delta s]^3} > 0,$$

$$g_2'(s_2) = \frac{2r\sigma^2[r\sigma^2 + (\underline{\theta}+2)\Delta s][1+2\underline{\theta}]}{[2r\sigma^2 + 3\Delta s]^3} > 0,$$

where I used the fact that $\overline{\theta} = \underline{\theta} + 1$. Both derivatives are positive, since $\underline{\theta} > 0$, $\Delta s \geq 0$, $r > 0$, $\sigma^2 > 0$. Therefore, equilibrium profit variances increase in own quality. ■

Proposition 2. *When quality is costless, Firm 2 will always choose the highest quality, $\overline{s}$.*

Proof. I will show that $dV_2/ds_2 > 0$ regardless of s_1. To simplify the expression for equilibrium utility, write the negative exponent as a product of two functions, f_2 and g_2:

$$V_2 = -\exp\left\{-\underbrace{\left(r\Delta s + \frac{1}{2}r^2\sigma^2\right)}_{f_2}\underbrace{\left[\frac{r\sigma^2 + (2\overline{\theta} - \underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2}_{g_2}\right\} = -e^{-f_2 g_2}.$$

We already found that $g_2' > 0$, and $f_2' = r$. Therefore,

$$\frac{dV_2}{ds_2} = -e^{-f_2 g_2}[-(f_2' g_2 + f_2 g_2')] = e^{-f_2 g_2}[f_2' g_2 + f_2 g_2'] = e^{-f_2 g_2}[r g_2 + f_2 g_2'].$$

Since both $f_2 > 0$, $g_2 > 0$, and, on the other hand, $g_2' > 0$, it follows that always $\frac{dV_2}{ds_2} > 0$, so the solution for s_2 must be at the right corner $s_2^* = \overline{s}$. ■

Proposition 3. *When quality is costless, and given that the second firm's choice is $s = \overline{s}$,*

(a) When consumer tastes are high enough ($\underline{\theta}$ exceeds some θ_*), Firm 1 will also choose $s = \overline{s}$, resulting in a ***non-differentiated*** *equilibrium $(s_1^*, s_2^*) = (\overline{s}, \overline{s})$ regardless of the values of r, σ^2.*

(b) Moreover, such a critical value θ_* always exists.

(c) When $\underline{\theta}$ falls below some θ_{**}, Firm 1 will choose $s_1^* = \underline{s}$, resulting in ***fully differentiated*** *equilibrium $(s_1^*, s_2^*) = (\underline{s}, \overline{s})$. Moreover, such a $\theta_{**} > 0$ always exists as long as $r\sigma^2 > 3$.*

Proof. (a) As before, write Firm 1's equilibrium utility in simpler form as

$$V_1 = -\exp\left\{-\left(r\Delta s + \frac{1}{2}r^2\sigma^2\right)\left[\frac{r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s}{2r\sigma^2 + 3\Delta s}\right]^2\right\} = -e^{-f_1 g_1}.$$

Then,

$$\frac{dV_1}{ds_1} = e^{-f_1 g_1}[f_1' g_1 + f_1 g_1'], = e^{-f_1 g_1}[-r g_1 + f_1 g_1'],$$

since $f_1' = -r$. Notice that the derivative's sign depends on the sign of $[-r g_1 + f_1 g_1']$, where the first term is negative while the second is positive (both $f_1 > 0$ and $g_1' > 0$, as shown in the proof to Proposition 2). Which term dominates will turn out to depend on the value of $\underline{\theta}$. To avoid dealing with the (strictly positive) denominators of g_1 and its derivatives, which do not change the sign, I bring the expression $[-r g_1 + f_1 g_1']$ to a common denominator and then focus on the numerator, so that

$$\begin{aligned}\operatorname{sgn}[-r g_1 + f_1 g_1'] = \operatorname{sgn}\{&-r[r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s]^2(2r\sigma^2 + 3\Delta s) + \\ &+2r\sigma^2\left(r\Delta s + r\sigma^2/2\right)[r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s][3 - 2(\overline{\theta} - 2\underline{\theta})]\},\end{aligned}$$

which, keeping in mind that $[r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s] > 0$ and $[3 - 2(\overline{\theta} - 2\underline{\theta})] = 1 + 2\underline{\theta}$, simplifies to

$$\begin{aligned}\operatorname{sgn}[-r g_1 + f_1 g_1'] = \operatorname{sgn}\{&-(r\sigma^2 + (1 - \underline{\theta})\Delta s)(2r\sigma^2 + 3\Delta s) + \\ &+2\sigma^2[r\Delta s + r^2\sigma^2/2](1 + 2\underline{\theta})\}.\end{aligned}$$

First I will prove that there exists a critical value θ_* above which $[-r g_1 + f_1 g_1'] < 0$. For this to happen, it must be the case that

$$2\sigma^2[r\Delta s + r^2\sigma^2/2](1 + 2\underline{\theta}) > (r\sigma^2 + (1 - \underline{\theta})\Delta s)(2r\sigma^2 + 3\Delta s),$$

or, after reduction,

$$r\sigma^2[3\Delta s + r\sigma^2] > [\, 6r\sigma^2\Delta s + 2r2\sigma^4 + 3(\Delta s)^2\,](1 - \underline{\theta}).$$

To find a θ_* high enough to work independent of Δs, I substitute "the worst possible values" $\Delta s = 0$ in the left-hand side and $\Delta s = 1$ in the right-hand side, resulting in

$$r^2\sigma^4 > (6r\sigma^2 + 2r^2\sigma^4 + 3)(1 - \underline{\theta}).$$

Therefore, when

$$\underline{\theta} > \underbrace{1 - \frac{r^2\sigma^4}{6r\sigma^2 + 2r^2\sigma^4 + 3}}_{\theta_*} \quad \Rightarrow \frac{dV_1}{ds_1} > 0.$$

This implies that the optimal choice of s_1 is at the right corner of the interval $[\underline{s}, \overline{s}]$ and that therefore $s_1^* = \overline{s}$ as claimed.

(b) Notice that since $0 < \frac{r^2\sigma^4}{6r\sigma^2+2r^2\sigma^4+3} < \frac{1}{2}$, it follows that $1 > \theta_* > 1/2$, which means that θ_* is always within the allowed range (0,1) for $\underline{\theta}$.

(c) To reverse the sign of dV_1/ds_1, we need the opposite inequality to hold:

$$(1 - \underline{\theta})[\, 6r\sigma^2\Delta s + 2r^2\sigma^4 + 3(\Delta s)^2 \,] > r\sigma^2[3\Delta s + r\sigma^2],$$

which reduces to

$$\underline{\theta} < 1 - \frac{r\sigma^2[3\Delta s + r\sigma^2]}{6r\sigma^2\Delta s + 2r^2\sigma^4 + 3(\Delta s)^2}.$$

To find a θ_{**} that works for all Δs, I substitute "the worst possible cases" $\Delta s = 1$ in the numerator and $\Delta s = 0$ in the denominator, yielding

$$\theta_{**} = 1 - \frac{r\sigma^2[3 + r\sigma^2]}{2r^2\sigma^4} = \frac{r\sigma^2 - 3}{2r\sigma^2},$$

which is smaller than 1 and positive as long as $r\sigma^2 > 3$. ■

Proposition 4. *When quality costs for both firms are given by $C(s) = bs^2$,*

*(a) If $\underline{\theta} < 1/2$ and $b > \frac{5}{16\underline{s}}$, both firms arrive at a non-differentiated equilibrium with **minimum** quality.*

*(b) If $\underline{\theta} > \frac{1}{2} + 8b\overline{s}$ and $b < \frac{1}{16\overline{s}}$, both firms arrive at a non-differentiated equilibrium with **maximum** quality.*

Proof. (a) Since we look for zero-differentiation equilibria at the corners, evaluate the quality derivatives of V_1 and V_2 at $\Delta s = 0$, looking for a relationship between them and costs that will make them positive or negative:

$$\left.\frac{dV_1}{ds_1}\right|_{\Delta s=0} = re^{-f_1 g_1 + rbs_1}\left[-\frac{1}{8} + \frac{1}{4}\underline{\theta} - 2bs_1\right]$$

$$\left.\frac{dV_2}{ds_2}\right|_{\Delta s=0} = re^{-f_2 g_2 + rbs_2}\left[\frac{3}{8} + \frac{1}{4}\underline{\theta} - 2b\underline{s}_2\right].$$

First notice that if $\underline{\theta} < 1/2$, then $\frac{dV_1}{ds_1}|_{\Delta s=0} < 0$, so Firm 1 will have an incentive to go for the minimum quality. For this to be an equilibrium, however, Firm 2 will also need to have incentives to stay there. V_2's derivative shows that

$$\left.\frac{dV_2}{ds_2}\right|_{\Delta s=0} < 0 \quad \text{whenever} \quad \underline{\theta} < 8b\underline{s} - 3/2.$$

The inequality on the right-hand side will always be satisfied when $8b\underline{s} - 3/2 > 1$ (because $\underline{\theta} < 1$ by assumption). So a sufficient condition for $\frac{dV_2}{ds_2}|_{\Delta s=0} < 0$ is

$$b > \frac{5}{16\underline{s}}.$$

When $\underline{\theta} < \frac{1}{2}$ and $b > \frac{5}{16\underline{s}}$, both firms' quality choice ($\underline{s}$) is a best response to the competitor's quality, so the quality pair $(s_1^*, s_2^*) = (\underline{s}, \underline{s})$ is a Nash equilibrium.

(b) Let us now determine when non-differentiation is possible at the maximum quality. Begin with Firm 1. If $\underline{\theta} > 1/2 + 8b\overline{s}$, then $dV_1/ds_1|_{\Delta s=0} > 0$, but $\underline{\theta}$ also needs to be less than 1. A necessary condition for this is $b < \frac{1}{16\overline{s}}$.

For Firm 2, the derivative evaluated at $\Delta s = 0$ will be positive as long as

$$\underline{\theta} > 8b\overline{s} - \frac{3}{2},$$

which, given that $\underline{\theta} < 1$, is possible only when

$$8b\overline{s} - \frac{3}{2} < 1 \Rightarrow b < \frac{5}{16\overline{s}}.$$

Collecting the conditions for both firms to go to $\overline{s}$, we get

$$\underline{\theta} > \frac{1}{2} + 8b\overline{s} \qquad b < \frac{1}{16\overline{s}},$$
$$\underline{\theta} > 8b\overline{s} - \frac{3}{2} \qquad b < \frac{5}{16\overline{s}}.$$

Taking their intersection provides the sufficient conditions

$$\underline{\theta} > \frac{1}{2} + 8b\overline{s} \qquad b < \frac{5}{16\overline{s}},$$

which guarantee that both firms choose $\overline{s}$. Since each firm is maximizing utility given the competitor's action, again $(s_1^*, s_2^*) = (\overline{s}, \overline{s})$ is a Nash equilibrium. ■

Proposition 5. *Let the firms' utility function be given by* $u(x) = \frac{x^{1-R}}{1-R}$ *with* $R = 2$, *and costs be uniformly distributed on an interval* $[\underline{c}, \overline{c}]$ *with mean* μ *and variance* σ^2, *with remaining assumptions unchanged. Then, if quality costs for both firms are given by the strictly increasing function* $C(s)$, *firms arrive at a* ***non-differentiated*** *equilibrium at the* ***minimum*** *quality.*

Proof. The utility function is

$$u(x) = \left.\frac{x^{1-R}}{1-R}\right|_{R=2} = -x^{-1},$$

and costs are uniformly distributed on an interval $[\underline{c}, \overline{c}]$ with mean μ and variance σ^2, so that $\underline{c} = \mu - \sqrt{3}\sigma$ and $\overline{c} = \mu + \sqrt{3}\sigma$ (by properties of the uniform distribution). Therefore, profits, as a function of cost, are also uniformly distributed, with lower bound $\underline{\Pi} = \Pi(\overline{c})$ and upper bound $\overline{\Pi} = \Pi(\underline{c})$. Since profits become negative when the realized cost exceeds the price, while the CRRA utility is defined only over positive wealth $x > 0$, we introduce a calibration constant $K > 0$ assumed to be large enough to ensure that $K + \underline{\Pi} > 0$ so that the utility function is well-defined. Since costs are bounded, for simplicity prices are also restricted[13] on $[\underline{c}, \overline{c}]$. The uniformly distributed profit has a probability density of $1/[\overline{c} - \underline{c}] = 1/2\sqrt{3}$ at every point in $[\underline{c}, \overline{c}]$. Therefore, expected utility is given by

$$\mathbb{E}u(x) = \int_{K+\underline{\Pi}}^{K+\overline{\Pi}} -x^{-1} \frac{1}{2\sqrt{3}\sigma} dx = \frac{-1}{2\sqrt{3}\sigma} \ln(x) \Big|_{K+\underline{\Pi}}^{K+\overline{\Pi}} = \frac{-1}{2\sqrt{3}\sigma} \ln\left(\frac{K+\overline{\Pi}}{K+\underline{\Pi}}\right),$$

where $\overline{\Pi}$ and $\underline{\Pi}$ are specific for each firm. For instance, for Firm 1,

$$\overline{\Pi}_1 \equiv (p_1 - \underline{c}) \left[\frac{p_2 - p_1 - \underline{\Delta}}{\Delta s}\right]; \qquad \underline{\Pi}_1 \equiv (p_1 - \overline{c}) \left[\frac{p_2 - p_1 - \underline{\Delta}}{\Delta s}\right].$$

[13] This restriction is realistic due to the existence of state coverage minima and price caps imposed by regulators. For example, imposing a minimum coverage standard, as is typical in the United States, effectively puts a lower bound on prices, since in the long run the insurer will not sell below the actuarially fair price. Similarly, binding price caps exist in insurance markets; as an example, in 2011, Massachusetts insurance regulators refused to approve a health insurance premium raise proposed by the industry.

For simplicity, from here on denote the expression $\frac{K+\overline{\Pi}}{K+\underline{\Pi}}$ with Q, so that $\mathbb{E}u_1(x) = \frac{-1}{2\sqrt{3}\sigma}\ln(Q_1)$.

The firm chooses a price to maximize expected utility based on the derivative

$$\frac{d\mathbb{E}u_1}{dp_1} = \frac{-1}{2\sqrt{3}\sigma}\left[\frac{K+\underline{\Pi}}{K+\overline{\Pi}}\right]\frac{dQ}{dp_1}.$$

It can be shown that, for Firm 1,

$$\frac{dQ_1}{dp_1} = \frac{\overline{\Pi}'(\underline{\Pi}+K) - \underline{\Pi}'(\overline{\Pi}+K)}{(\underline{\Pi}+K)^2} < 0.$$

To see this, consider the expressions $\overline{\Pi}_1'$ and $\underline{\Pi}_1'$:

$$\overline{\Pi}_1' = \left[\frac{p_2 - p_1 - \underline{\Delta}}{\Delta s}\right]\frac{p_2 - p_1 - \underline{\Delta} - 2(p_1 - \underline{c})}{\Delta s},$$

$$\underline{\Pi}_1' = \left[\frac{p_2 - p_1 - \underline{\Delta}}{\Delta s}\right]\frac{p_2 - p_1 - \underline{\Delta} - 2(p_1 - \overline{c})}{\Delta s}.$$

Since $\overline{c} > \underline{c}$, from these expressions it follows that $\overline{\Pi}_1' < \underline{\Pi}_1$. Taking into account that, by definition, $\overline{\Pi} > \underline{\Pi}$, it follows that the expression

$$\overline{\Pi}'(\underline{\Pi}+K) - \underline{\Pi}'(\overline{\Pi}+K) < 0,$$

and therefore $d\mathbb{E}u_1/dp_1 > 0$. This implies that Firm 1 will choose the highest price $p_1 = \overline{c}$ in the Bertrand equilibrium regardless of what Firm 2 does.

The same result can be shown for the Bertrand price charged by Firm 2. The upper and lower bound on profits for Firm 2 are

$$\overline{\Pi}_2 \equiv (p_1 - \underline{c})\left[\frac{\overline{\Delta} - p_2 + p_1}{\Delta s}\right]; \qquad \underline{\Pi}_2 \equiv (p_1 - \overline{c})\left[\frac{\overline{\Delta} - p_2 + p_1}{\Delta s}\right].$$

Likewise, I proceed to show that $\frac{dQ_2}{dp_2} < 0$. For Firm 2, similarly $\overline{\Pi}_2' < \underline{\Pi}_2'$, because

$$\overline{\Pi}_2' = \left[\frac{\overline{\Delta} - p_2 + p_1}{\Delta s}\right]\frac{\overline{\Delta} - p_2 + p_1 - 2(p_2 - \underline{c})}{\Delta s},$$

$$\overline{\Pi}_2' = \left[\frac{\overline{\Delta} - p_2 + p_1}{\Delta s}\right]\frac{\overline{\Delta} - p_2 + p_1 - 2(p_2 - \overline{c})}{\Delta s},$$

and $\overline{c} > \underline{c}$, therefore rendering $\overline{\Pi}'(\underline{\Pi} + K) - \underline{\Pi}'(\overline{\Pi} + K) < 0$, from which it follows that $dQ_2/dp_2 < 0$ and therefore $d\mathbb{E}u_2/dp_2 > 0$. Hence, in the Bertrand equilibrium Firm 2 will also charge the highest price $p_2 = \overline{c}$, regardless of the action of Firm 1.

The Bertrand equilibrium in the game's second stage is therefore $p_1 = \overline{c}$ and $p_2 = \overline{c}$.

Now suppose quality s costs $C(s)$, where C is an increasing function of s. The equilibrium utility attained by each firm is then

$$V_1^*(s_1) = \frac{-1}{\sqrt{3}\sigma} \ln\left(\frac{K + 2\sqrt{3}\sigma\underline{\theta}^2}{K}\right) - C(s_1); \quad V_2^*(s_2) = \frac{-1}{\sqrt{3}\sigma} \ln\left(\frac{K + 2\sqrt{3}\sigma\overline{\theta}^2}{K}\right) - C(s_2).$$

Clearly, V_1^* and V_2^* are maximized when $C(s)$ is minimized and $s_1 = s_2 = \underline{s}$. It is optimal for firms not to differentiate by choosing the lowest possible quality but extracting the highest possible price. This example confirms the intuition that oligopolists can use their market power to force a low-quality outcome as a way of economizing on costs.

The above example with risk aversion $R = 2$ results in corner price solutions, but this need not always be the case for CARA utility. For example, if one considers $R = 3$, it is possible to show that the Bertrand prices are interior on the interval $(\underline{c}, \overline{c})$, but they cannot be obtained analytically as functions of s. ■

8.2 Asymptotics of Baseline Model at $\Delta s \to 0$

Lemma. *As $\Delta s \to 0$, equilibrium demand, profits and their variances converge to the following quantities:*

$$\lim_{\Delta s \to 0} D_1^* = \lim_{\Delta s \to 0} D_2^* = \frac{1}{2}$$
$$\lim_{\Delta s \to 0} \mathrm{Var}(\Pi_1^*) = \lim_{\Delta s \to 0} \mathrm{Var}(\Pi_2^*) = \frac{1}{4}\sigma^2$$
$$\lim_{\Delta s \to 0} \Pi_1^* = \lim_{\Delta s \to 0} \Pi_2^* = \frac{r\sigma^2}{4}.$$

Proof. It is sufficient to find the limit of the expression $\frac{p_2^* - p_1^*}{\Delta s}$, which enters D^*, Π^* and $\mathrm{Var}(\Pi^*)$. Recall that equilibrium prices p_1^* and p_2^* satisfy the system of reaction functions

$$p_1(p_2) = \frac{\mu_c \Delta s + (r\sigma^2 + \Delta s)(p_2 - \underline{\Delta})}{r\sigma^2 + 2\Delta s},$$

$$p_2(p_1) = \frac{\mu_c \Delta s + (r\sigma^2 + \Delta s)(p_1 + \overline{\Delta})}{r\sigma^2 + 2\Delta s}.$$

By subtracting the first equation from the second, obtain

$$\frac{p_2^* - p_1^*}{\Delta s} = \frac{r\sigma^2 + \Delta s}{2r\sigma^2 + 3\Delta s}(\overline{\theta} + \underline{\theta}).$$

In the limit $\Delta s \to 0$, it becomes

$$\lim_{\Delta s \to 0} \left[\frac{r\sigma^2 + \Delta s}{2r\sigma^2 + 3\Delta s}(\overline{\theta} + \underline{\theta}) \right] = \frac{(\overline{\theta} + \underline{\theta})}{2} = \underline{\theta} + \frac{1}{2}.$$

$$\text{Since } D_1^* = \frac{p_2^* - p_1^*}{\Delta s} - \underline{\theta}, \quad \Rightarrow \lim_{\Delta s \to 0} D_1^* = 1/2.$$

$$\text{Analogously, } D_2^* = \overline{\theta} - \frac{p_2^* - p_1^*}{\Delta s} \Rightarrow \lim_{\Delta s \to 0} D_2^* = \underbrace{\overline{\theta} - \underline{\theta}}_{1} - \frac{1}{2} = \frac{1}{2}.$$

From here it easily follows that $\lim_{\Delta s \to 0} \mathrm{Var}(\Pi_1^*) = \lim_{\Delta s \to 0} \mathrm{Var}(\Pi_2^*) = \frac{1}{4}\sigma^2$. To prove the result for profits, observe that the equilibrium markups converge to

$$\lim_{\Delta s \to 0} m_1^* = \lim_{\Delta s \to 0} \left[\frac{(r\sigma^2 + \Delta s)[r\sigma^2 + (\overline{\theta} - 2\underline{\theta})\Delta s]}{2r\sigma^2 + 3\Delta s} \right] = \frac{r\sigma^2}{2},$$

$$\lim_{\Delta s \to 0} m_2^* = \lim_{\Delta s \to 0} \left[\frac{(r\sigma^2 + \Delta s)[r\sigma^2 + (2\overline{\theta} - \underline{\theta})\Delta s]}{2r\sigma^2 + 3\Delta s} \right] = \frac{r\sigma^2}{2},$$

so it is enough to multiply the limit of the markup times that of demand. This shows that, as differentiation approaches zero, the model converges to classical Bertrand competition with a homogeneous good where firms do not differentiate, charge the same price and split the market equally. ■

References

[1] Bester, H., "Quality uncertainty mitigates product differentiation," *RAND Journal of Economics* 29 (1998), 828-844.

[2] Blazenko, G., "The Design of an Optimal Insurance Policy: Note." *American Economic Review*, 75(1) (1985), 253–255.

[3] Cho, I. and D. Kreps, "Signalling Games and Stable Equilibria," *Quarterly Journal of Economics* 102 (1987), 179-221.

[4] Crampes, C. and A. Hollander, "Duopoly and quality standards," *European Economic Review*, 39(1) (1995), 71–82.

[5] Cummins, D. and B. Venard, *Handbook of international insurance: between global dynamics and local contingencies.* (New York: Springer, 2007).

[6] Ecchia, G. and L. Lambertini, "Minimum Quality Standards and Collusion," *Journal of Industrial Economics* 45(1) (1997), 101–113.

[7] Eliashberg, J. and R. Winkler, "Risk Sharing and Group Decision Making." *Management Science*, 27(1) (1981), 1221–1235.

[8] Freund, R., "The Introduction of Risk into a Programming Model," *Econometrica* 24(3) (1956), 253–263.

[9] Froot, K., *The Financing of Catastrophe Risk.* (Cambridge, MA: National Bureau of Economic Research, 1999).

[10] Froot, K., D. Scharfstein and J. Stein, "Risk Management: Coordinating Corporate Investment and Financing Policies," *Journal of Finance* 48 (1993), 1629–1658.

[11] Gabszewicz, J. and J.-F. Thisse, "Price competition, quality and income disparities," *Journal of Economic Theory* 20(3) (1979), 340–359.

[12] Garella, P. G. and E. Petrakis, "Minimum quality standards and consumers' information," *Economic Theory*, 36(2) (2008), 283–302.

[13] Hardelin, J. and S. Lemoyne de Forges, "Raising Capital in an Insurance Oligopoly Market," *Geneva Risk and Insurance Review* 37 (2012), 83–108.

[14] Harrington, S., *Insurance Deregulation and the Public Interest.* (Washington, D.C.: The AEI Press, 2000).

[15] Hellwig, M., "A Sequential Approach to Modeling Competition in Markets with Adverse Selection." Mimeo, University of Bonn, (1986).

[16] Hellwig, M., "Some Recent Developments in the Theory of Competition in Markets with Adverse Selection," *European Economic Review* 31 (1987), 319–325.

[17] Maxwell, J. W., "Minimum quality standards as a barrier to innovation," *Economics Letters* 58(3) (1998), 355–360.

[18] Mayers, D. and C. Smith, "On Corporate Demand for Insurance," *Journal of Business* 55 (1982), 281–296.

[19] Mussa, M. and S. Rosen, "Monopoly and Product Quality," *Journal of Economic Theory* 18 (1978), 301–317.

[20] OECD, *Policy Issues in Insurance, 9: Terrorism Risk Insurance in OECD Countries.* (OECD Publishing: Paris, 2005).

[21] Polborn, M., "A model of an oligopoly in an insurance market," *Geneva Papers on Risk and Insurance Theory* 23 (1998), 41–48.

[22] Raviv, A., "The Design of an Optimal Insurance Policy," *American Economic Review*, 69(1) (1979), 84–96.

[23] Raykov, R., "Catastrophe Insurance Equilibrium with Correlated Claims," *Theory and Decision*, forthcoming, (2014).

[24] Rees, R., H. Gravelle and A. Wambach, "Regulation of Insurance Markets," *Geneva Papers on Risk and Insurance Theory* 24 (1999), 55–68.

[25] Rothschild, M. and J. Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 90 (1976), 629-649.

[26] Sandmo, A., "On the theory of the competitive firm under price uncertainty," *American Economic Review* 61 (1971), 65–73.

[27] Schmalensee, R. and R. D. Willig, *Handbook of Industrial Organization*, Vol. 1. (North-Holland: New York, 1989).

[28] Shaked, A. and J. Sutton, "Relaxing price competition through product differentiation," *Review of Economic Studies* 49 (1982), 3–13.

[29] Smith, C. and R. Stulz, "The Determinants of Firm's Hedging Policies," *Journal of Financial and Quantitative Analysis* 20 (1985), 391–405.

[30] Stulz, R., "Optimal Hedging Policies," *Journal of Financial and Quantitative Analysis* 19 (1984), 127–140.

[31] Tirole, J., *The Theory of Industrial Organization*. (MIT Press: Cambridge, MA, 1988).

[32] Toshimitsu, T. and N. Jinji, "Minimum quality standards under asymmetric duopoly with endogenous quality ordering: a note," *Journal of Regulatory Economics* 26(2) (2004), 189–199.

[33] Valletti, T., "Minimum quality standards under Cournot competition," *Journal of Regulatory Economics*, 18(3) (2000), 235–245.

[34] Wambach, A., "Bertrand competition under cost uncertainty," *International Journal of Industrial Organization* 17(1999), 941-951.